Helios™ Context Continuity Architecture (HCCA™)

Whitepaper v1.0 — Codex-Centric Memory and Rules Continuity

Author / Originator: Anthony

Brand: Helios™

Version: v1.0 (Formatting Revision A)

Date: 2025

Author’s Note

This architecture emerged from extended, hands-on experimentation with large language models across long conversations and separate sessions, where continuity of memory, rules, and user preferences proved essential but unreliable when left to platform-managed mechanisms. Modern tools, including AI systems, were used to assist in drafting and refinement; however, the architectural decisions, scope boundaries, and claims presented here reflect direct human design choices and empirical observation. Responsibility for all claims, limitations, and interpretations rests solely with the author.

Abstract

Large language models (LLMs) exhibit fundamental continuity limitations: information loss across long conversations due to finite context windows and loss of state between separate conversation threads. Vendor-managed memory features, where available, are typically opaque, non-portable, and platform-dependent. This paper introduces Helios™ Context Continuity Architecture (HCCA™) v1.0, a codex-centric approach that demonstrates how memory and rules, including user preferences, can persist across context windows and across separate conversation threads through the deliberate creation and reuse of Codexes—structured artifacts that externalize saved data and rulesets. HCCA™ v1.0 establishes a minimal, explicit, and proven architecture for continuity that is independent of vendor-managed memory systems.

1. Introduction

Large language models are fundamentally stateless systems. Any apparent memory is typically confined to the active context window and is lost when that window is exceeded or when a new conversation thread is created. This behavior undermines long-running projects, stable system identity, and consistent application of user preferences and rules.

2. Problem Statement

HCCA™ v1.0 addresses context window loss, thread boundary loss, and reliance on opaque vendor-managed memory systems. These limitations prevent reliable continuity and reproducibility in long-running AI-assisted work.

3. Scope and Claims (v1.0)

This whitepaper claims only three capabilities: memory and rules continuity across context windows; memory and rules continuity across separate conversation threads; and the creation and reuse of Codexes as structured artifacts containing saved data and rulesets.

4. Definitions

Memory refers to persisted information required for continuity, including facts, definitions, and working state. Rules refer to constraints and preferences governing behavior, including user preferences. A Codex is a structured, persistent artifact that stores memory and rulesets in explicit form.

5. Architectural Overview

HCCA™ v1.0 is codex-centric. Continuity is achieved by externalizing memory and rules into Codexes that exist independently of any single conversation and can be reintroduced when continuity is required.

6. Continuity Across Context Windows

Finite context windows cause truncation of earlier information. By externalizing critical memory and rules into Codexes, HCCA™ v1.0 enables continuity across long conversations.

7. Continuity Across Separate Conversation Threads

New conversation threads normally begin without access to prior memory or rules. Reintroducing the Codex restores continuity without relying on platform-managed memory.

8. Codex Creation as the Core Mechanism

Codex creation is the central contribution of HCCA™ v1.0. By treating memory and rules as explicit artifacts, continuity becomes portable, inspectable, and user-controlled.

9. Limitations

Reintroduction of Codexes may be manual. Behavioral variance across models may occur. No claims are made regarding automation, enforcement, or determinism.

10. Related Work

HCCA™ differs from vendor-managed memory systems and generic retrieval approaches by focusing on explicit artifact reuse rather than hidden or incidental persistence.

11. Authorship, Naming, and Use

Helios™ Context Continuity Architecture (HCCA™) v1.0 is a named architecture originated by Anthony. This document establishes authorship, scope of claims, and historical priority.

12. Conclusion

HCCA™ v1.0 demonstrates that memory and rules continuity—including user preferences—can persist across context windows and separate conversation threads through the creation and reuse of Codexes. This establishes a minimal, proven foundation for AI continuity.